

Eliciting naturalistic cortical responses with a sensory prosthesis via optimized microstimulation

John S Choi¹, Austin J Brockmeier², David B McNiel¹, Lee M von Kraus¹,
José C Príncipe² and Joseph T Francis^{1,3,4}

¹ Joint Program in Biomedical Engineering, Polytechnic of Institute of NYU and SUNY Downstate, Brooklyn, NY, USA

² Department of Electrical and Computer Engineering, University of Florida, Gainesville, FL, USA

³ Department of Physiology and Pharmacology, State University of New York Downstate Medical School, Robert F Furchgott Center for Neural and Behavioral Science, Brooklyn, NY, USA

⁴ Department of Biomedical Engineering, University of Houston, Houston, TX, USA

E-mail: john.choi@downstate.edu

Received 23 January 2016, revised 10 May 2016

Accepted for publication 29 June 2016

Published 12 August 2016



Abstract

Objective. Lost sensations, such as touch, could one day be restored by electrical stimulation along the sensory neural pathways. Such stimulation, when informed by electronic sensors, could provide naturalistic cutaneous and proprioceptive feedback to the user. Perceptually, microstimulation of somatosensory brain regions produces localized, modality-specific sensations, and several spatiotemporal parameters have been studied for their discernibility. However, systematic methods for encoding a wide array of naturally occurring stimuli into biomimetic percepts via multi-channel microstimulation are lacking. More specifically, generating spatiotemporal patterns for explicitly evoking naturalistic neural activation has not yet been explored. **Approach.** We address this problem by first modeling the dynamical input–output relationship between multichannel microstimulation and downstream neural responses, and then optimizing the input pattern to reproduce naturally occurring touch responses as closely as possible. **Main results.** Here we show that such optimization produces responses in the S1 cortex of the anesthetized rat that are highly similar to natural, tactile-stimulus-evoked counterparts. Furthermore, information on both pressure and location of the touch stimulus was found to be highly preserved. **Significance.** Our results suggest that the currently presented stimulus optimization approach holds great promise for restoring naturalistic levels of sensation.

Keywords: microstimulation, optimal control, sensory prostheses

(Some figures may appear in colour only in the online journal)

1. Introduction

Loss of somatosensation could one day be treated by direct electrical stimulation of the central nervous system.

Assessment of the naturalness of microstimulation-induced sensations is difficult in animal models, since it has been largely achieved by training animals to report these sensations as being conceptually ‘higher’ or ‘lower’ in some regard than natural stimuli [1–4]. For example, [4] demonstrated that a static nonlinear transformation of touch pressure could match detection and relative amplitude discrimination rates seen in natural touch experiments. The subjects could similarly report spatial comparisons (i.e. ‘more medial than’) as well. Other



Original content from this work may be used under the terms of the [Creative Commons Attribution 3.0 licence](https://creativecommons.org/licenses/by/3.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

work has explored sensitivities to variations of microstimulation temporal pattern, spatial variation, and level of randomness [5–8].

While these psychophysical studies show the potential for discriminability of such induced sensations, it is unlikely that the simple heuristically chosen pulse patterns, often involving only a single electrode at a time, used in these studies are sufficient to recreate natural cortical activation and natural sensations. Indeed, in humans, constant-amplitude pulse trains applied to single electrodes in the ventral caudal thalamus evoke percepts that are both place and modality-specific, and yet ‘unnatural’ in feeling [9, 10]. Although such signals could provide useful feedback in a brain machine interface, developing more biomimetic spatiotemporal patterns remains an open problem. As these patterns increase in complexity, the difficulty of assessing performance psychophysically in animals increases, perhaps prohibitively.

A likely way to increase the realism of microstimulation-induced sensations is to use dynamic, biomimetic encoding algorithms that are specific to a subject, the implanted electrodes, and the state of the neuronal circuit. A relatively simple method is to, on every electrode, inject current pulses according to the predicted naturally occurring firing rate. This method was employed by [11] to replicate the spiking activity of a damaged hippocampal region. Reference [12] showed that this method could modulate cortical activity in a naturalistic fashion when applied to electrodes implanted in the dorsal root ganglion. Unfortunately, these methods are predicated on the assumption of a one-to-one correspondence between a stimulation pulse and elicited spike. Electrophysiological evidence in [13] suggests that each pulse instead produces a spatiotemporal blur of activity involving many cells. Indeed, for the range of currents likely to be useful for evoking percepts, a single microelectrode would have direct effects on neuronal elements within 30–100 μm of their conducting areas [14, 15].

Given these effects, a more reasonable approach is to deliver microstimulation pulse patterns at locations upstream of the target population in a manner that trans-synaptically induces desired downstream activation. To this end, we have developed a model-based control method capable of eliciting naturalistic responses in the somatosensory cortex via optimized patterns of intra-thalamic microstimulation (ITMS). These patterns spatiotemporally resembled naturalistic spike rates, and their evoked responses preserved most of the information of touch parameters.

2. Methods

In this study, two separate microelectrode arrays (see figures 2(a)–(b)) were implanted in order to record and stimulate synchronously. The first, situated in the forelimb representation of the VPL thalamus, delivered the microstimulation, and the second, situated in the corresponding

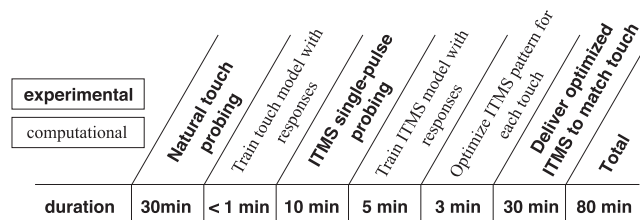


Figure 1. Experimental timeline. ITMS: intra-thalamic microstimulation.

projection area in S1, measured ongoing neural activity during stimuli (natural touch or microstimulation).

We investigated the following procedure in rats (see figure 1). A set of downstream responses to rectangular skin indentations of varying pressure, duration, and location are obtained to serve as templates. Probing ITMS is then delivered and the neural responses are used to train a linear state-space model of the effects of VPL microstimulation. A controller then optimizes a set of pulse patterns that, in terms of the model, approximates the natural downstream responses as closely as possible. The optimal patterns are then applied in VPL, the responses are recorded, and the similarity of the responses are assessed. Herein, we consider multi-electrode recordings of the local field potentials (LFP). As an alternative, sets of spike trains or spike counts can be used [16–19], but as a continuous signal the LFP is simpler for state-space modeling. In addition, we can use mean-square error and correlation as metrics amenable to highly efficient convex optimization. This study concentrates on characterizing the neural responses to both natural taction and optimized microstimulation along with their similarity. Studying the performance enhancement it offers behaviorally is left as future work.

2.1. Surgical methods

All animal procedures were approved by the SUNY Downstate Medical Center Institutional Animal Care and Use Committee. Nine female Long-Evans rats (250–350 g) were acutely implanted with electrode arrays in VPL and S1 (see figure 2(b)) under urethane anesthesia. In the first six animals, the microelectrode array (MicroProbes Inc.) in VPL was a 2×8 grid of 70% platinum 30% iridium 75 μm diameter microelectrodes, with 500 μm between the rows, and 250 μm inter-electrode spacing within the rows. The shank lengths were custom designed to fit the contour of the rat VPL [20]. Both rows were identical and the shaft lengths for each row, from medial to lateral, were {8, 8, 8, 8, 8, 7.8, 7.6, 7.4} mm. In the remaining animals, we used a 32 channel 4-shank multi-contact silicon array (NeuroNexus A4x8-10mm-200-500-703-CM32) in VPL. This array was used in place of the traditional microwire array due to its higher contact density in VPL and lower insertion force. Each shank contained eight contacts separated by 200 μm . In our insertions we found that 2–4 shanks picked up spike responses to touch on a subset of contacts, which corresponded with known maps of rat VPL.

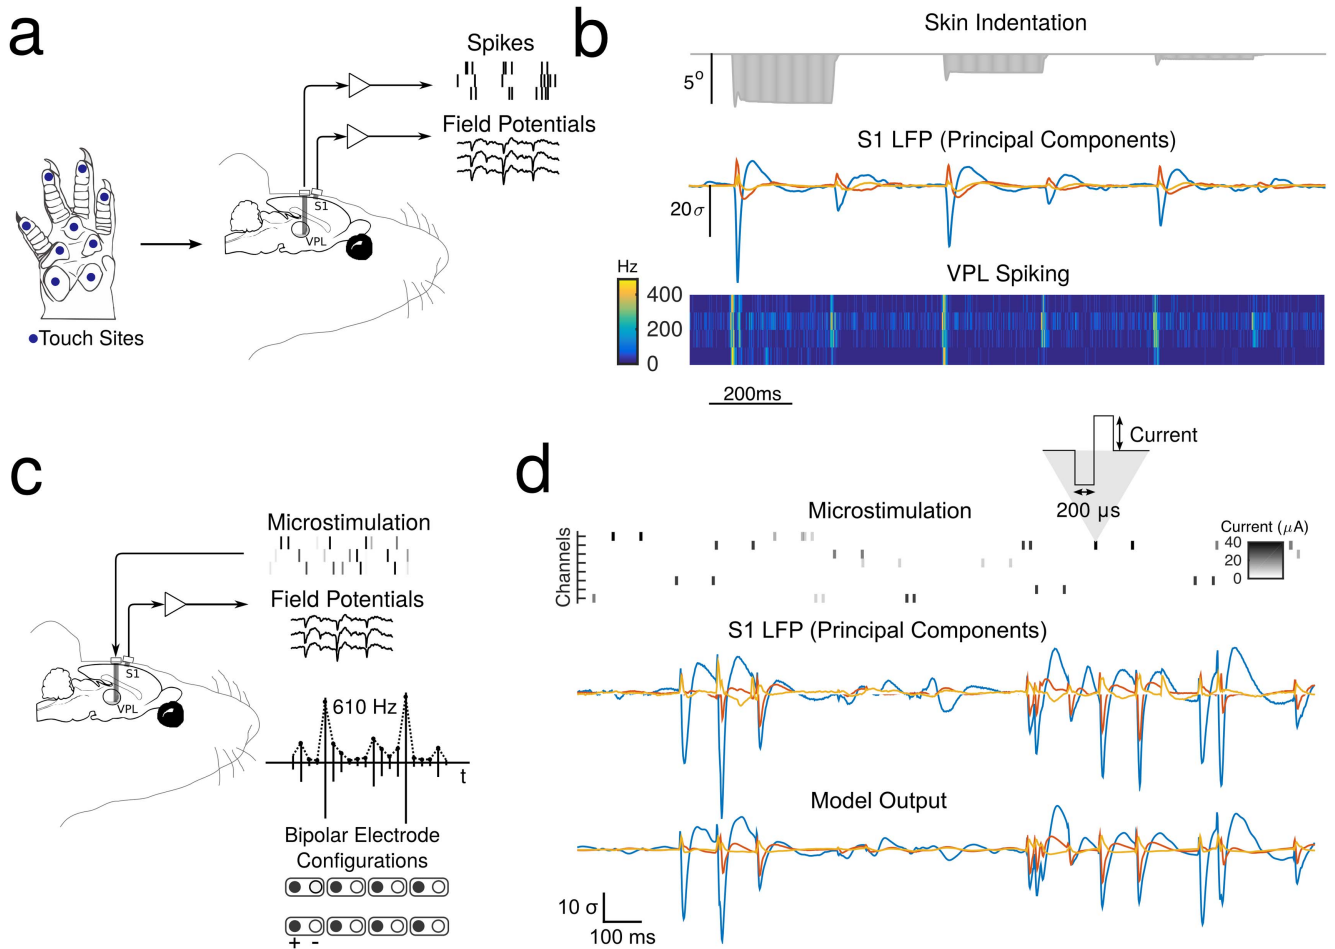


Figure 2. (a) Natural touch stimulation delivered to various touch sites on the ventral forepaw. (b) Touch responses measured by field potentials in S1 and spike rate in VPL. (c) Multichannel microstimulation delivered through the VPL array, with LFP recorded in S1. (Inset) Parameterization of microstimulation patterns: For each channel, the pulse train's amplitude is modulated by an envelope signal. Each stimulation channel corresponds to an adjacent pairs of electrodes. (d) Sample microstimulation modeling sequence (top), corresponding S1 response trajectory (middle), and linear model output (bottom).

In the first three animals, the cortical electrode array (Blackrock Microsystems) was a 32 channel Utah array (see figure 2(b)). The electrodes are arranged in a 6×6 grid excluding the four corners, and each electrode is 1.5 mm long, with a spacing of $400 \mu\text{m}$. In the remaining six animals, we instead used a 4-shank silicon multi-contact array (NeuroNexus A4x8-5mm-100-400-703-CM32). This array allowed us to measure activity along the dorsoventral axis, a spatial dimension that is fixed in the Utah array. This enabled us to test our optimization not only across the surface of cortex, but also different cortical layers. Electrode arrays were positioned using stereotaxic coordinates for the digit region of S1 (4.0 mm lateral and 0.5 mm anterior to bregma) [21, 22]. The VPL electrode array was centered on stereotaxic coordinates for the hand representation in the medial subdivision of VPL [20].

In order to recover stable neural activity following array insertion, we waited a period of 2 h post-insertion before starting neural recording and stimulation, and a stable plane of anesthesia was maintained through small supplemental doses of urethane.

2.2. Recording neural responses during natural touch probing

Multichannel LFP were collected during tactile stimulation with a neural signal acquisition system (RZ2, Tucker-Davis Technologies). Broadband field potentials were band-pass filtered with cutoff frequencies at 5 and 200 Hz and sampled at 610 Hz. This filtered signal is what we refer to in this work as LFP. The average post-touch-onset LFP responses for each (touch site, indentation, and hold time) combination were collected up to 300 ms following touch onset. These responses were used as target waveforms for optimizing the multichannel intrathalamic microstimulation (ITMS). Some example waveforms and a channel map capturing the spatial extent of excitation are shown in figures 4 and 6.

The single unit activity (SUA) in VPL was recorded using standard spike-sorting techniques. Peri-stimulus time histograms (PSTH) were taken for each distinct natural stimulus condition, with a bin size equal to the sampling period at which the LFP was sampled, i.e., 1.63 ms within a window of time from the moment of touch onset to 50 ms after touch offset. The single-unit activity in VPL was collected for two reasons: firstly, to compare optimized thalamic

microstimulation to native VPL spiking; and secondly, to measure the natural response reproduction accuracy of a microstimulation signal whose amplitude follows the same rate modulation as VPL spikes—similar to a method used in [23].

Physical touch stimulation was administered to 3–9 sites on the ventral surface of the right forepaw with a precision tactor. The touch patterns consisted of a touch-hold-release sequence parameterized by a pressure and duration. For each site, three different skin indentations and two different touch durations {150, 250} ms were applied. A shuffled series of 25 instances of each of the six patterns were presented with random time intervals in between the touches. The pressures were chosen to evoke responses ranging from threshold to near-saturation: indentation depths: {0.025, 0.2, 0.6} mm. Due to the small size of the digits, we did not attempt indentations larger than 0.6 mm.

The tactile probe was driven by a DC motor (Maxon RE-25) fitted with an optical encoder (Maxon HEDS-55). Timing of contact and level of skin indentation were controlled via the motor's angle, which was accomplished by a proportional derivative controller implemented onboard the neural recording system. The endpoint of the probe was a circular shaft 1 mm in diameter. For the first three animals, the probe was attached at the end of a beam 9 cm in length attached directly to the motor. In the remaining six animals in this study, the probe was attached to a mechanical slide via a rack and pinion mechanism (gear diameter = 0.438"). Precise control of the mechanism was achieved via control of the motor's rotation, which translated into linear motion of the probe used to touch discrete locations on the forepaw. The amplitude of the force applied by the bar is directly proportional to the skin indentation by Hooke's Law. For the purposes of this work, however, we present results in terms of the lever angle in degrees.

2.3. Recording and modeling LFP response to electrical stimulation

To train a model of the cortical LFP response to VPL microstimulation input, we used a random multichannel pulse sequence. A multichannel microstimulator (IZ2, Tucker-Davis Technologies) delivered the pulses to non-overlapping bipolar pairs of adjacent electrodes spanning the VPL array (see figures 2(c) and (d) for illustration). Eight bipolar configurations were used for 6 of 9 animals where 16-channel arrays were implanted, and 16 bipolar configurations in the remaining animals, where 32 channel arrays were used. We used bipolar configurations since they produced less stimulation artifact in cortical recording channels than monopolar ones. Pulses were symmetric and biphasic (200 μ s per phase). Symmetric pulses were chosen because of their relative safety compared to asymmetric pulses [24–26]. The choice of pulse width was chosen on the basis of consistency with similar studies [13, 27]. Three different stimulation amplitudes {10, 20, 30} μ A were used in the probing phase of our work for the first three animals. However, on the remaining six animals, we stimulated with a more comprehensive range of

amplitudes: {7, 12, 20, 30, 40} μ A, which evoked responses ranging from subthreshold to saturation when using bipolar configurations. The procedure for pulse timing and delivery is as follows: each pulse-to-pulse interval was drawn from an exponential distribution, and a stimulating configuration and amplitude was chosen uniformly at random from all configurations and predesignated current amplitudes. This input distribution is a multi-input variant of the random-amplitude Poisson stimulation train, which has been used with success in the past for modeling neuronal responses [28–30]. In 6 of 9 animals, the probing distribution was interleaved with pulse doublets where pairs of pulses were delivered with inter-pulse-intervals chosen from a discrete list of intervals {20, 50, 75, 100, 200} ms. The use of pulse doublets was meant to enable a set of analyses separate from the results of this study. The mean frequency of pulse delivery was varied from subject to subject but ranged from 3 to 8 Hz in the first three animals and 12–18 Hz in the remaining six. The total duration of the probing sequences also varied from 6 to 18 min and each unique combination of configuration/amplitude was repeated 50–240 times.

To facilitate the modeling and control steps, we represent this pulse sequence as a discrete-time amplitude envelope that modulates a constant-frequency pulse train (see figure 2(c)). The frequency was set equal to the sampling rate of LFP recording (610 Hz). At each time step, we treat each channel as the amplitude gain on a 1 μ A pulse through each unique stimulation configuration (adjacent bipolar pair). In this study, this resulted in 8 or 16 distinct input channels depending on the electrode array used. By using bipolar configurations, stimulation artifacts were generally small and brief. The unfiltered waveforms never exceeded 0.5 mV, well below the clipping level of the amplifiers, and lasted for less than 200 μ s after the end of a stimulation pulse. Since all signals were initially filtered at broadband (0.2–8.5 kHz) frequencies and sampled at 24.4 kHz, filter ringing was also minimal. This allowed us to use a simple sample-and-hold method for blanking a 480 μ s period of time starting from the start of a pulse. This blanking was applied digitally to the broadband 24.4 kHz signal prior to further processing.

Several signals of interest can be obtained from extracellular recording within cortex following microstimulation of upstream thalamic nuclei, including SUA and LFP, among others [31]. While each of these cortical signals is a potential target for control using thalamic microstimulation, our group ultimately decided to use cortical LFP in this work. We chose LFP as opposed to SUA for its comparative efficacy in decoding touch parameters [32, 33] in a similar implant setting. It has also been shown to be a robust signal for decoding motor activity in a variety of contexts [34–36]. LFP has gained some recent interest for brain machine interfaces due to its relative robustness over long periods of implantation compared to SUA. While the origin of LFP remains controversial, it is believed that LFP, being a measure of aggregate extracellular voltage resulting from membrane currents, represents dendritic input within a small (<200 μ m) region surrounding the electrode. We also noticed that the stimulation artifact, although small in amplitude after the

aforementioned blanking, was difficult to completely eliminate from being misclassified as valid multi-unit spiking due to the large variety of artifact waveforms (multiple configurations, amplitudes) and their similarity in some cases to spike waveforms. Additionally, the continuous nature of the LFP signal allows for the application of more straight-forward comparison than the discrete-valued spiking signal. For LFP, distance measures such as the traditional mean-squared error and cross-correlation can be used, whereas distances between spike trains are more varied and involve tuning parameters.

2.3.1. Model. A discrete-time linear state-space model was trained using the subspace identification algorithm [37, 38]. Subspace system identification methods estimate system parameters by finding, using a dataset of input–output pairs, a projection that maps past inputs and outputs to future outputs. From this projection, a low-dimensional sequence of state variables can be extracted, along with parameters that describe their associated temporal dynamics and relations to the observed output. A description of the model follows, but we refer to [37] for algorithmic details.

Let $\mathbf{u}(t) \in \mathbb{R}^m$, $\mathbf{y}(t) \in \mathbb{R}^p$ denote the input (multichannel microstimulation amplitude envelope) and output (neural readout), respectively, at time t . Let $\mathbf{x}(t) \in \mathbb{R}^n$ denote the state vector at time t . The state-space equations relating the current state with the state at time $t + 1$ and the output at time t are

$$\mathbf{x}(t + 1) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t) + \epsilon_x, \quad (1)$$

$$\mathbf{y}(t) = \mathbf{C}\mathbf{x}(t) + \epsilon_y, \quad (2)$$

where $\epsilon_x \sim N(\mathbf{0}, \mathbf{Q})$ and $\epsilon_y \sim N(\mathbf{0}, \mathbf{R})$ are Gaussian white (uncorrelated in time) disturbances. The system parameters are hence the matrices $(\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{Q}, \mathbf{R})$ which are estimated using a subspace method (algorithm 4.8 in [37]).

In preliminary experiments, we sometimes found that this optimization procedure would ‘undershoot’ the amplitude of microstimuli necessary to reproduce the desired response. This was due to the linear model not capturing stimulation thresholds, and the the controller, as a result, relying on sub-threshold amplitude regimes for control. To solve this problem, we modeled a microstimulation threshold using a ‘gate’ function that attenuated the input if it was below a certain threshold value and left it unchanged otherwise. Precisely, for each input channel i

$$\text{gate}(u_i) = \begin{cases} u_i & \text{if } u_i \geq \text{threshold} \\ au_i & \text{otherwise,} \end{cases} \quad (3)$$

where $a \in (0, 1]$ is an attenuation factor used to pre-scale subthreshold input values before being passed into the state space equation (1).

In our experiments, the model dimensions and gate parameters were set as follows: an 8 or 16 dimensional input for representing distinct microstimulation configurations was used. In 3 of 9 animals, the output was simply the recorded LFP on 32 channels, but in 6 of 9 animals, we used a reduced representation consisting of the first 12–15 principal components that captured 99.7% of the observed instantaneous variance in the output. A target state dimensionality of 50 was

chosen empirically as a tradeoff between model complexity and prediction accuracy. We set the gate attenuation factor and the threshold value based on a separate probing pulse sequence that used a lower set of current amplitudes. We found that suitable threshold values ranged from 4 to 10 μA . The attenuation factor was difficult to measure directly due to the wide variability and nonlinearity of the responses near threshold. We instead set this factor by hand to 0.1 or 0.2. We noticed that very small values of the attenuation factor (values approaching 0) were not used because of their deleterious effect during stimulus optimization.

2.4. Optimizing microstimulation patterns

Our goal is to optimize multichannel ITMS control inputs in order to reproduce, as closely as possible, the natural touch responses for each unique touch type presented. Specifically, we penalize the deviation between our system output and some desired template neural response trajectory over a finite time period. We first pose this optimization problem as a quadratic cost function with linear equality and inequality constraints. This type of problem, known as a quadratic program, is well-studied and specialized algorithms exist for efficient polynomial-time solution [39]. One strategy for solving this type of control problem is to solve for not only the optimal control input over the time horizon, but also the state variables. However, the states are implicitly a function of the input (1), and the optimization procedure must enforce this dynamical relationship through linear constraints. This formulation, although containing more variables and constraints, is actually beneficial computationally when problem structure is exploited [40].

The optimization problem also contains inequality constraints due to our representation of electrical amplitude. Since the probing microstimulation sequence only consisted of pulses of a single polarity, we choose to constrain the input to a single polarity to keep all optimized stimulations in an approximately linear regime. Changing the polarity switches the order in which current is sunk or sourced on each of the two adjacent electrodes, and response to negative polarities are not linearly related to the response to positive polarities [41]. For this reason, we impose a non-negativity constraint on all inputs as well as a maximum input bound to keep solutions within the range of current amplitudes explored during probing. As in the modeling step, the pulse sequence consisted of a 610 Hz train for each stimulating channel.

Let $\mathbf{y}^d(t)$ denote the desired neural trajectory at time t . We assume that at time t , a desired neural trajectory \mathbf{y}^d for the times $t + 1, \dots, t + T - 1$ is available. The horizon T governs the amount of future time that the controller considers in optimizing control inputs. In practical (causal) applications, this desired signal could be the output of a predictive response model that, using sensor information available up to t , outputs a predicted neural response for $T - 1$ future time points. Alternatively, \mathbf{y}^d could also be a precomputed/recorded neural trajectory. For this study, we treated \mathbf{y}^d as completely known, i.e., we set it to the peristimulus trial average of the natural response for each touch condition.

The primary goal of the controller is to minimize the distance between $\mathbf{y}(t)$, the system output under the applied input sequence, and the desired signal $\mathbf{y}^d(t)$. We define the quadratic cost at stage τ as

$$\ell(t, \mathbf{x}(\tau)) = \|\mathbf{y}^d(t) - \mathbf{C}\mathbf{x}(\tau)\|^2. \quad (4)$$

The optimization goal is to minimize, within the aforementioned constraints, this cost over T time steps and can be stated as

$$\begin{aligned} & \underset{\mathbf{u}, \mathbf{x}}{\text{minimize}} \quad \sum_{\tau=t}^{t+T} \ell(t, \mathbf{x}(\tau)) \\ & \text{subject to } 0 \leq \mathbf{u}(\tau) \leq I_{\max}, \\ & \mathbf{x}(\tau+1) = \mathbf{A}\mathbf{x}(\tau) + \mathbf{B}\mathbf{u}(\tau), \tau = t, \dots, t+T-1, \end{aligned} \quad (5)$$

where the optimization is performed over values of $\mathbf{u}(t), \dots, \mathbf{u}(t+T-1)$ and $\mathbf{x}(t+1), \dots, \mathbf{x}(t+T)$, and I_{\max} is the maximum current limit, which we set to the largest current value used during the probing sequence. The model's dynamics are enforced by equality constraints relating the current input and state ($\mathbf{x}(\tau), \mathbf{u}(\tau)$) to the next state $\mathbf{x}(\tau+1)$. The system evolution in this case is deterministic and the optimization does not depend on the density of ϵ_x or ϵ_y in equation (1).

We include a secondary objective that penalizes large currents. This is accomplished by adding a term $\mu\|\mathbf{u}(t)\|^2$ to the cost function in equation (4) that penalizes the squared norm of the input, where μ is a weighting parameter that controls the relative importance of this penalty. Similarly, one could economize on the amplitude of a low-pass filtered version of the input by adding $\lambda v(t)^2$ to the stage cost where $v(t+1) = (1-\alpha)v(t) + \alpha \sum_i u_i(t)$ —a rudimentary single-pole low-pass filter with a pole at $(1-\alpha)$. This has the effect of penalizing high amplitude, slowly varying input patterns. We noticed that without this penalty, some of the less effectual inputs would be driven to continuously stimulate at significant amplitudes. Although these inputs would be accomplishing nominally better output tracking over the control horizons, they were stimulating at significant supra-threshold amplitudes without much pertinent effect on the measured field potentials. We suspect that these were channels that influenced a part of S1 that was only partially picked up by our recording array. In our experiments, the relative weighting factors μ and λ were hand-chosen for each animal according to a tradeoff between current injection and tracking error. In contrast, the low-pass filter parameter α was fixed to $1/(\tau_{lp}F_s + 1)$ where F_s is the sampling frequency and τ_{lp} , the filter time constant, was set to 100 ms.

This method of penalizing a filtered version of the inputs is very similar to the method used by [18]. In that work, however, the optimization was done over raw current waveforms, and the penalty on slow current injection served primarily the purpose of limiting charge build-up, which is well known for causing electrode corrosion or tissue damage near contacts. In our work, charge balance is immediately restored with each biphasic pulse, since our optimization is on the amplitude envelope of a stereotyped pulse train. Instead, the

filtered penalty puts a selective cost on slow, sustained trains of pulses.

The optimal control problem in equation (5) is quadratic in the control inputs and states. In our formulation, the variable being optimized is a concatenation $\mathbf{z} = (\mathbf{u}(t), \mathbf{x}(t+1), \mathbf{u}(t+1), \dots, \mathbf{u}(t+T-1), \mathbf{x}(t+T))$. This leads to a set of equality constraints enforcing the system dynamics of inputs and states in adjacent points in time, and a set of inequality constraints enforcing control input bounds. Since equation (5) is convex (i.e., its 2nd derivative with respect to the \mathbf{z} is positive definite for all \mathbf{z}) and its equality and inequality constraints are linear, this problem can be solved tractably by convex optimization methods [39]. By exploiting the structure induced by the equality constraints only being enforced for adjacent time points, the running time is $\mathcal{O}(T(n+m)^3)$ —a vast improvement over the running time if the structure was not exploited $\mathcal{O}(T^3(n+m)^3)$. We refer the reader to [40] for details on the specific algorithm used to solve (5) via an interior point method.

The input gate feature introduced to our model makes the system nonlinear. The state transition with an input gate is $\mathbf{x}(t+1) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{gate}(\mathbf{u}(t))$, where $\mathbf{gate}(\mathbf{u})_i = \mathbf{gate}(u_i)$. To deal with this, on each iteration, the system at each time t is linearized around the current value of the input, denoted $\tilde{\mathbf{u}}(t)$, and hence the nonlinear input-dependence can be replaced by a linear time-varying term. Precisely, the gated input in the state transition equation can be replaced by its first-order Taylor approximation

$$\begin{aligned} \mathbf{gate}(\mathbf{u}) &\approx \mathbf{gate}(\tilde{\mathbf{u}}) + \frac{\partial}{\partial \mathbf{u}} \mathbf{gate}(\tilde{\mathbf{u}})(\mathbf{u} - \tilde{\mathbf{u}}) \\ &= \text{diag}(\mathbf{g}(\tilde{\mathbf{u}}))\mathbf{u}, \\ \mathbf{g}(\tilde{\mathbf{u}})_i &= \begin{cases} 1, & \text{if } \tilde{u}_i \geq \text{threshold} \\ a, & \text{otherwise.} \end{cases} \end{aligned}$$

Hence the state transition equation can be treated as a time-varying, but linear, function $\mathbf{x}(t+1) = \mathbf{A}\mathbf{x}(t) + \hat{\mathbf{B}}(t)\mathbf{u}(t)$ where $\hat{\mathbf{B}}(t) = \mathbf{B}\text{diag}(\mathbf{g}(\tilde{\mathbf{u}}(t)))$. As a damping measure, a combination of new and current solutions is taken after each iteration as $\mathbf{z} := \beta\mathbf{z}_{\text{new}} + (1-\beta)\mathbf{z}$. In our experiments, we found that setting β equal to 1.0 initially and scaling β by 0.97 on each iteration until $\beta = 0.3$ was suitable for finding good solutions.

Another method we explored was the iterative linear quadratic regulator (iLQR) [42]. This also used successive time-varying linearizations of the system dynamics in order to cope with the gate nonlinearity. iLQR also uses a different method for updating the input and states based on LQR, which produces linear state-feedback policies of the form $\mathbf{u}(t) = \phi_t(\mathbf{x})$ instead of just input trajectories. We found that in our implementations, iLQR's input solutions did not differ significantly with those of the interior point method.

Initially, the optimal control input was found for each touch condition (site, amplitude, duration) offline. The desired trajectory in each case was the trial-averaged natural touch response with $\tau = 0$ corresponding to touch onset and $T = \text{hold duration} + 50$ ms. The microstimulation pattern begins at touch onset and continues up until T .

Once found, the optimized ITMS patterns were applied through the VPL array. The patterns for each touch type were applied in the same order and timing as the original natural touch stimuli for each forepaw location. We define the term *virtual touch* to refer to stimulating with optimized patterns of microstimulation corresponding to a particular type of natural touch.

2.5. Comparison with rate-based stimulation

An alternative method for encoding microstimulation patterns is to have them resemble how the stimulated neurons would normally spike in response to touch stimulation. This was the approach used in [23] where spikes were ‘played back’ on the stimulating array with the same spatiotemporal timing and with each spike represented by a stimulation pulse of an appropriate amplitude. We noticed that applying this directly to our experimental setting was problematic for two reasons. One was that many units in VPL showed a fairly high background firing rate, so playing back pulses unrelated to the stimuli would cause unintended activation of S1. Even when a wide range of substitution pulse amplitudes were exhaustively attempted, these patterns would not reliably evoke effective stimulus tuning. The other was that often two or more distinct units would be detected on each channel, and even more would be present in each stimulating configuration. Combining these units into a single stimulation signal required a way to deal with this redundancy.

To address these issues, we first extracted the PSTHs for each detected VPL unit for each touch condition. The PSTH was calculated so as to represent, for each bin, the expected (average) number of spikes relative to stimulus onset. Precisely, for the i th bin relative to touch onset, $\text{PSTH}_i = \text{count}_i / N_{\text{trials}}$ where count_i is the total number of spikes to occur in that bin for all N_{trials} trials. Then, the background rate (total number of spikes divided by the duration of the recording) was subtracted from each PSTH. This step was designed to mitigate the effect of background firing. Then, for each stimulating configuration (bipolar configurations on adjacent electrodes), the unit with the highest peak in the background-subtracted PSTH was selected as the representative unit for the stimulus configuration. This was meant to select the unit with the most unambiguous touch-stimulus-tuning. Then, a global gain was applied to the representative PSTHs to produce a stimulation signal in μA . Since the PSTHs were scaled such that each bin’s value represented the average number of spikes in that bin, the gain is precisely the μA per spike. For each animal, the optimal gain was determined by searching exhaustively in increments of 1 from 5 to 15 for the most accurate S1 LFP reproduction in the touch site with the largest response. This gain was then fixed and the rate-based stimulations were applied for the remaining touch sites. This method was applied in 3 of 9 animals in our study.

2.6. Decoding touch parameters from responses

Given a single multichannel response, how accurately can the touch parameters of location, amplitude, and duration be decoded? Can virtual touch responses, having been optimized for naturalness, be decoded with similar levels of accuracy? To measure this discriminability, we performed a set of classification experiments in which the touch condition (location, amplitude, duration) was predicted from the peristimulus touch response.

This was first done separately for virtual and natural touch, to see how much information the neural responses in each modality provides about the touch parameters. Ideally, however, virtual responses to different touch parameters should not only be discriminable from each other, but should be well-separated along the same boundaries as natural responses. To test this, we attempted to classify virtual touch responses under a single ‘generalized’ classifier whose feature space is defined using both virtual and natural responses, but whose classification boundaries are based solely on the natural responses. We will first describe the algorithm used in these experiments and present classification rates in the next section.

The decoding procedure, in both the individual and generalized classifiers, consists of supervised dimensionality reduction followed by nearest-mean classification. We assume the multichannel LFP responses (within the post-onset window of T samples), given the label, can be treated as random vectors and modeled using a unimodal distribution such as a multivariate normal. Under this assumption, each sample can be assigned the label of the nearest peristimulus average. To take into account covariance, linear discriminant analysis (LDA) is used to project the responses into a lower dimensional subspace. LDA has a closed form solution given by a generalized eigenvalue problem defined by the between-class and within-class covariance matrices [43]. Since the responses have many dimensions ($p \cdot T$), principal component analysis (PCA) is performed before computing the covariance matrices. Then each response is projected into a reduced subspace before being assigned the label of the nearest peristimulus average. In the results that follow, we chose the reduced dimensionality by cross validation, but it is at most one less than the number of classes [43]. For most of the results we use a post-onset window of 300 ms, which corresponds to $T = 367$ samples at the raw sampling frequency of 1220 Hz. For short duration touches we also analyze the effect of window length on classification performance, this is done by computing a different LDA projection for every window length.

The individually trained classifiers used responses from exclusively natural or virtual touch when learning the LDA projection and when assigning the final class label. In the generalized classifier, however, the LDA projection was performed using both types of responses, but the final classification was performed by assigning the label of the nearest natural-touch mean. All classifiers were validated through 8 Monte-Carlo divisions of data (2/3 train, 1/3 test) and the results are presented in the next section.

To directly measure how much information the decoded touch parameters provided about the true touch labels, the mutual information (in bits) between the two discrete label variables was computed empirically. As in the computation of classification rate, each unique combination of touch parameters was treated as a discrete label, and the expected conditional entropy of the decoded touch parameter label given the true label was subtracted from the marginal entropy of the touch label. To compute the information rate in bits s^{-1} , the information per touch was multiplied by the average rate of touch delivery for each rodent.

3. Results

3.1. Microstimulation model

The cortical LFP response to microstimulation was modeled with the linear state-space model described in the previous section. Figure 2(d) shows an example of the model output for a segment of the probing sequence. Across animals, the model on average accounted for $39.0\% \pm 16.8\%$ of the variance in the time periods that were within 400 ms after a stimulation pulse, and the average correlation coefficient was 0.61. Accuracy was highest during the intervals immediately following stimulation pulses, with the best performance in a window from 0 to 14.8 ± 2.0 ms following a pulse. As shown in figure 5(a), the model accuracy degraded gradually for longer time windows. No correlation of model accuracy with the overall LFP signal energy ($r = 0.061$, $p = 0.88$, $N = 9$) was detected, where energy was measured as the rms voltage values on all channels. Similarly, we did not detect a correlation with stimulation frequency, measured as pulses per second ($r = 0.4$, $p = 0.28$, $N = 9$). However, a correlation might have been present between the temporal current density, measured in $\mu A s^{-1}$ ($r = 0.62$, $p = 0.0739$, $N = 9$). Model error was not correlated with electrode implantation depth ($p = 0.751$).

3.2. Accuracy in reproducing natural responses

The optimized ITMS waveforms elicited neural responses that were spatiotemporally similar to their natural counterparts across touch sites and patterns. Across all conditions and rats, the correlation coefficient between the average natural and virtual responses was 0.78 ± 0.05 . If the comparison was made only for periods of time that were within 100 ms of touch onset, the correlation coefficient was 0.90 ± 0.03 . Figure 3 shows two segments of tactor position, natural LFP responses, optimized microstimulation, and the LFP responses to virtual touch. Trial-to-trial variability was similar on average between natural and virtual responses. The variability for each touch condition was measured by subtracting the trial average response and taking the root-mean-square (rms) value of the centered responses. The median rms variability across rodents was $79 \mu V$ and $86 \mu V$ for natural and virtual touch responses, respectively, but significance was not detected ($p = 0.1$, Wilcoxon signed rank test). In one isolated case (Rat

D in subsequent figures and tables), the variability was fairly low, especially for virtual touches, with an rms variability of $24 \mu V$ for virtual touch and $43 \mu V$ for natural touch.

Examples of the average temporal responses to different touch patterns are shown in figure 4. Natural touch elicits a strong, brief potential 9–15 ms after touch onset, followed by a recovery period lasting 150–200 ms. Another temporal feature is the smaller negative potential that occurs shortly after touch-offset when the actuator starts rising away from the touch site. The corresponding optimized microstimulation pattern is shown in figure 4(c). Figure 4(d) shows the resulting average LFP response.

Figures 6(a)–(b) show examples of the spatial responses from two different sites (digit 1, digit 4) on the rat hand. Displayed are the maximum negative deviations in the trial-averaged virtual and natural touches. Each channel of the S1 recording array is shown in its actual spatial arrangement. Natural taction of the sites d1 and d4 activated two overlapping but clearly distinct zones which were replicated to a high degree of accuracy ($r = 0.91 \pm 0.04$ in that particular animal). Overall, the spatial reproduction accuracy, measured by correlation coefficient, was $r = 0.72 \pm 0.22$ across all touch patterns in all animals. We found that the spatial reproduction accuracy with the 32 channel Utah array was 24.3% higher ($p = 0.02$) than that of the 32 channel Michigan probe. The Utah array, since it distributes its channels across the cortical surface rather than vertically like in the Michigan probe, could capture somatotopic variations more unambiguously while sacrificing information in the dorso-ventral (laminar) axis.

Figure 6(c) shows, for a representative animal, a comparison (natural versus optimized ITMS) of the energy output of S1, where we define the energy as the combined rms voltage of the multichannel response in the response window. Generally, the response energy for each natural touch type was well matched by its ITMS counterpart ($r = 0.81 \pm 0.13$). We also compared the experimental reproduction accuracy with the theoretical accuracy of the model output. The modeled reproduction accuracy was 5.8% higher than that achieved *in vivo*, and this relation was significant ($p < 0.001$, $N = 546$ touch conditions in all animals). Figure 6(d) shows the model versus *in vivo* reproduction accuracy for all touch conditions in all animals.

The control reproduction accuracy varied as a function of aspects of the touch stimuli and corresponding responses. In particular, we noticed that accuracy was better in cases where the target touch response was large compared to background noise. We measured this by computing the signal to noise ratio defined as: $SNR = 10 \log_{10}(\sum_i ||\mathbf{y}^d_i||^2) / (\sum_i ||\mathbf{y}^d_i - \mathbf{y}_i||^2)$, where \mathbf{y}_i is the observed response on trial i and \mathbf{y}^d is the trial average. Figure 5(b) shows overall that control reproduction accuracy (correlation coefficient) was correlated ($r = 0.601$, $p < 0.001$) with this SNR.

A two-way repeated measures ANOVA was conducted to assess the effect of touch strength and duration on control reproduction accuracy. Significant effects were found for both strength ($F = 6.7$, $p = 0.0017$) and duration ($F = 20.1$,

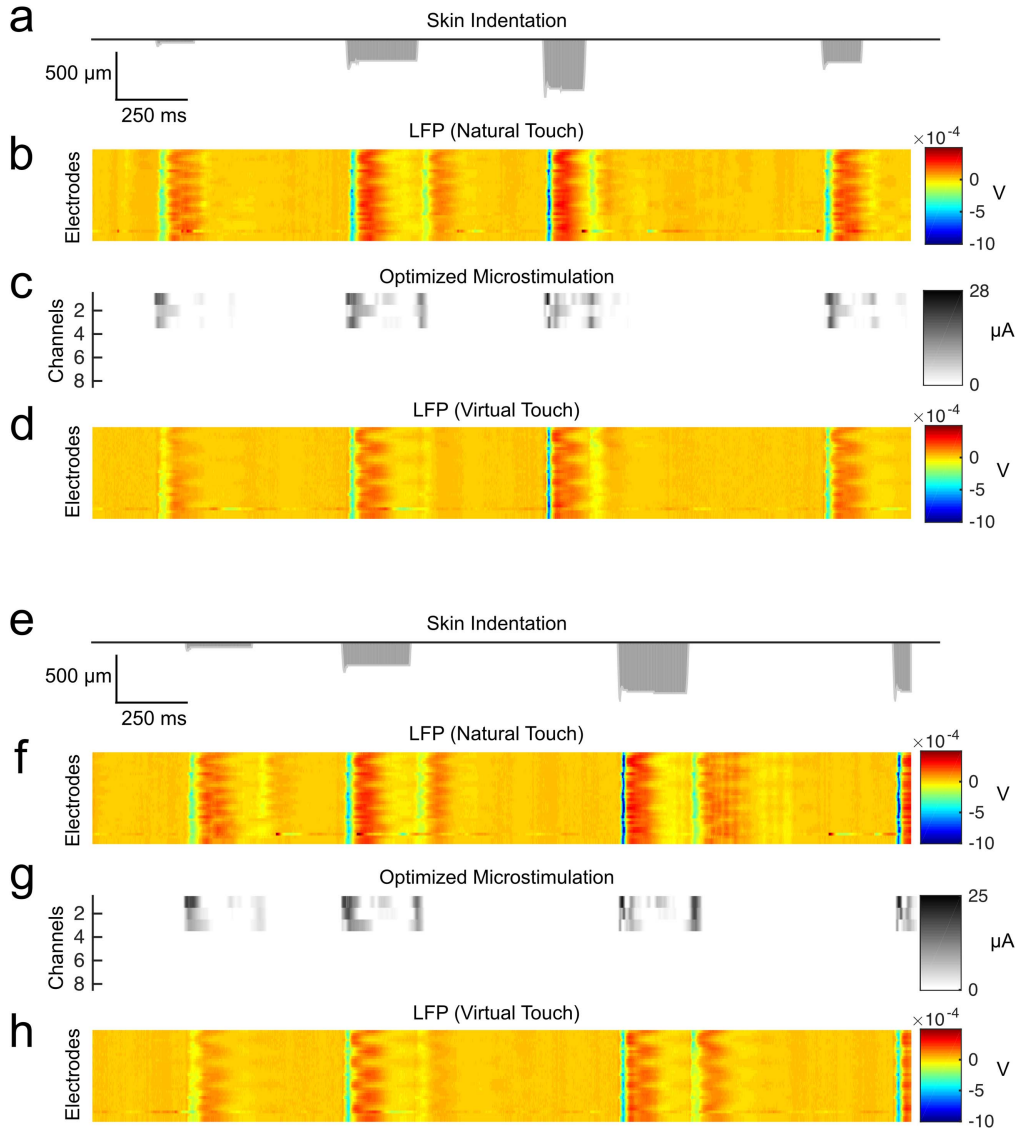


Figure 3. (a) Skin indentation measured by tactor position as a function of time for a sample of data (negative is towards the skin). The shaded portion of the curve represents periods of time in which the tactor was in contact with the skin. (b) Multichannel S1 LFP for 32 electrode channels during natural touch. Channels are sorted by their overall response amplitude for this touch site. (c) Optimized microstimulation delivered through eight channels. The shade indicates the amplitude of current, and channels are shown sorted by their overall usage for this touch site. (d) S1 LFP during optimized microstimulation, or ‘virtual touch.’ (e)–(h) Similar plots for a different dataset.

$p = 2.7 \times 10^{-5}$, $N = 73$ touch sites). Post-hoc tests were performed between strengths within each duration group (three comparisons per group) and between durations within each strength group (one comparison for each group) are shown in figure 5(c). Significance thresholds were Bonferroni corrected to $\alpha = 0.05/(\text{number of comparisons})$. No effect of the interaction between strength and duration was detected ($F = 1.4$, $p = 0.25$). A significant correlation between control accuracy and model accuracy was not detected ($r = 0.32$, $p = 0.40$).

The global accuracy scores for different touch patterns, averaged over touch sites, is shown in figure 5(c). Stronger touches were more easily reproducible than medium or light touches, and shorter-duration touch patterns had a higher accuracy than longer patterns. Figure 5(d) shows accuracies for all touch sites in all animals, sorted in increasing order within and among animals.

The accuracy of virtual touch responses was also measured in terms of their natural variability using Mahalanobis distances. For each touch condition, the distance was computed from the virtual touch response mean with respect to the mean and covariance of the natural response. This comparison was done in the same time window for which the stimulation was optimized, and because of the high dimensionality of the responses ($p \cdot T$), PCA was performed to reduce this dimensionality to $N_{\text{trials}} - 1$ components. The distance between the virtual touch mean, represented as a vector x of principal component scores, and the natural response distribution with mean μ and covariance Σ is

$$d(x, \mu, \Sigma) = \sqrt{(y - \mu)^T \Sigma^{-1} (y - \mu)}.$$

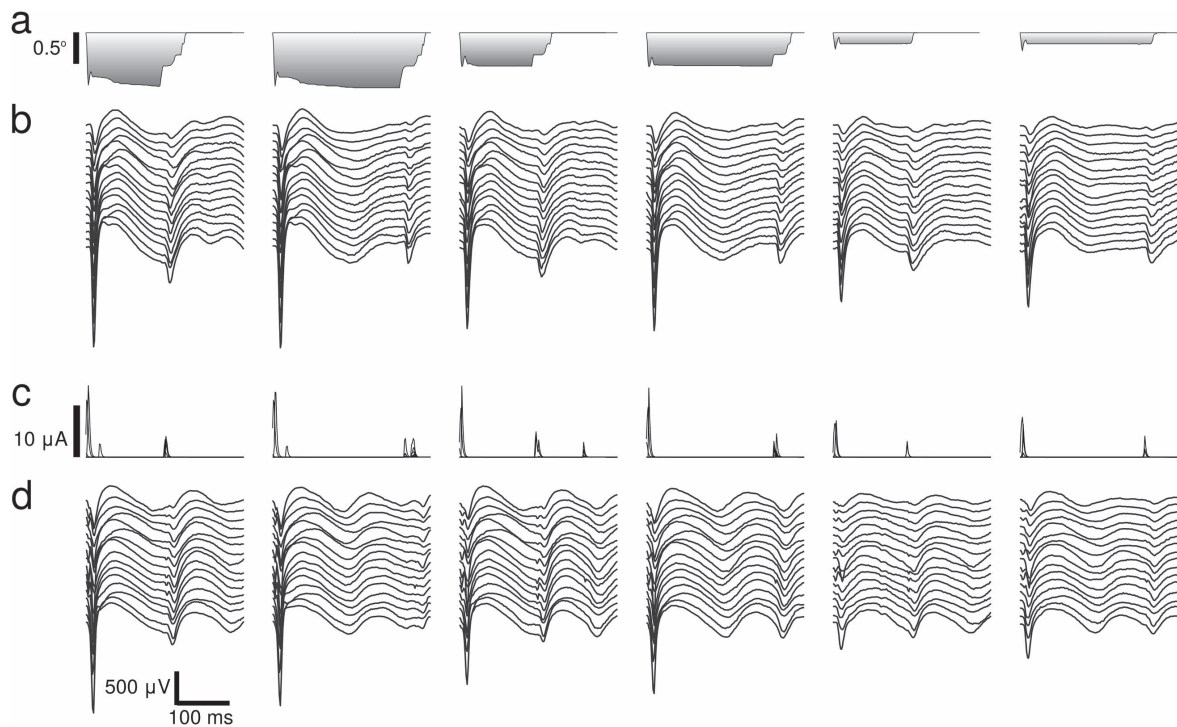


Figure 4. Post-touch-onset trace of the trial-average LFP for six touch patterns on a single touch site. (a) Trial-average factor angular position (negative is towards the skin). The shaded portion of the top plot corresponds to the period of time when the actuator was in contact with the skin. (b) Average multichannel LFP response to natural touch, each curve representing an S1 recording channel. Channels are sorted according to their overall response amplitude for this touch site. (c) Optimized microstimulation, each curve represents a distinct stimulation channel. (d) Average microstimulation LFP response.

We found that virtual touch responses showed specificity for their corresponding touch conditions. This was quantified by measuring, for each touch condition, the average distance over all conditions excluding the current one, and we found that the average ‘unmatched’ distance was 1.23-fold larger than the normal, matched distance ($p < 0.0001$ two-sided rank-sum test).

The performance of rate-based stimulation was quantified in the same way, and its distances were 1.38-fold larger than that of optimized virtual touch ($p < 0.0001$ two-sided rank-sum test). These distances were more variable across conditions. The standard deviation over distances was 2.3 for virtual touch ($N = 438$) and 4.2 for rate-based stimulation ($N = 27$). No significant difference in the median distances of the unmatched virtual touch and rate-based stimulation was detected ($p = 0.59$ two-sided rank sum test). A spatiotemporal comparison of rate and optimized microstimulation patterns will be given later in this section. Figure 5(e) shows a comparison of Mahalanobis distances for virtual touch, unmatched virtual touch, and rate-based stimulation.

3.3. Optimized pattern characteristics

We examined the output of our optimization procedure in terms of the timing of pulses and their spatiotemporal properties. Consistently with known somatotopy, VPL spatial

current injection for different touch sites followed a progression from medial to lateral as the touch site varied from medial to lateral on the hand. Any given virtual touch primarily used 1–3 bipolar configurations (2–6 stimulating electrodes), and the number of configurations used across all touch sites spanned 4–5 configurations. Most of the pulses occurred in a short burst from 4 to 8 ms after onset. This coincides with observations that the initial response latency to taction is ~ 9 ms, and the latency of VPL stimulation is ~ 2 ms.

One question of interest is how closely the optimized ITMS patterns resembled single-unit activity measured on the stimulating electrodes, and this was first quantified in a strictly spatial sense. For each touch condition, the sum was taken of total current injection on each stimulating channel. This was compared with the peri-stimulus spike count on each channel. Since these channels corresponded with bipolar electrode configurations, we calculated the spike count of the most touch-sensitive unit detected in each configuration in exactly the same way as used during rate-based stimulation. Figure 7(a) shows a representative example of the spatial variation of current injection and VPL spiking as a function of touch site. As the touch site is varied from one side of the hand to the other, the charge injected on each electrode, as well as native spiking, follows a spatial progression from medial to lateral on the VPL array. This spatial overlap was

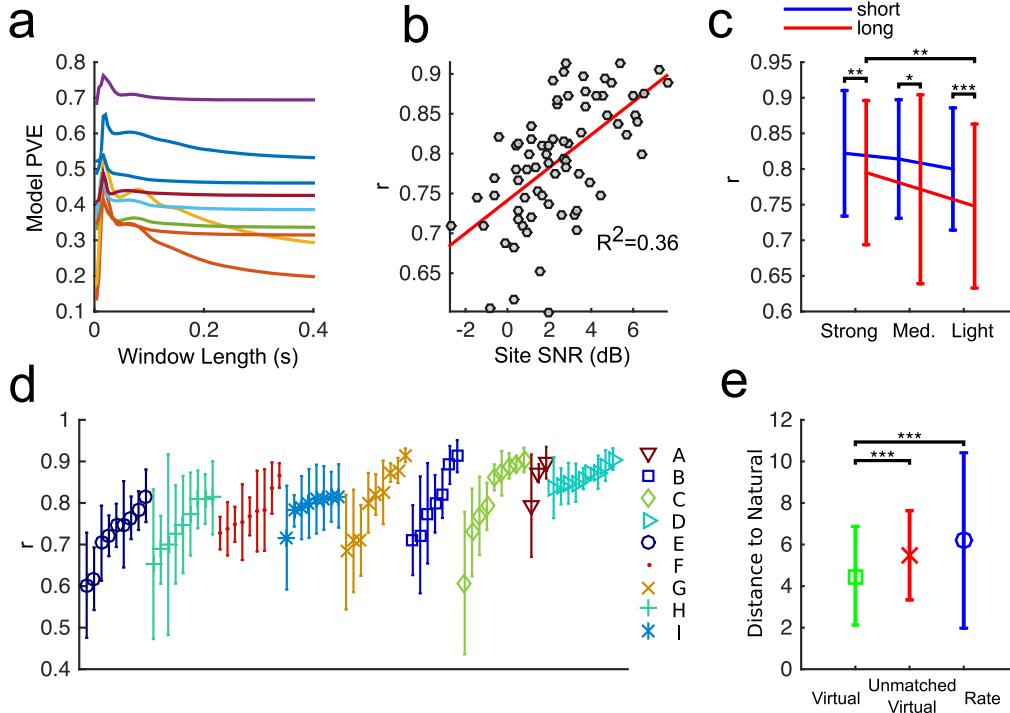


Figure 5. (a) Model accuracy when restricting comparison to variable-length time-windows following every stimulation pulse. Each curve represents one animal. PVE: percent variance explained. (b) Control reproduction accuracy as a function of the signal-to-noise ratio (SNR) of the natural touch LFP response. Each point represents a touch site. (c) Control reproduction accuracy as a function of touch parameters (*: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$, Bonferroni corrected. Error bars: ± 1 s.d. $N = 73$ touch sites, all animals). (d) Control reproduction accuracy of all touch sites in all animals. (e) Mahalanobis distances between the responses of natural touch and those of virtual touch ($N = 438$), unmatched-virtual touch ($N = 438$), and rate-based stimulation ($N = 162$). Significance and variability indicated the same way as in (c).

more concretely analyzed by taking the correlation coefficients between the spatial patterns of current injection and spiking across all touch conditions for each animal. The average correlation was $r = 0.46 \pm 0.38$, and 7 of 9 animals showed a significant correlation at a significance threshold of $\alpha = 0.05$, of which the correlation values ranged from $r = 0.26$ to 0.90 .

To quantify temporal correlation with VPL spike rate, the background-subtracted PSTHs (see methods) were compared with optimized ITMS for all animals. Figure 7(b) shows, for two different touch durations, the post-onset spike rate and the corresponding optimized ITMS. For each time point, the displayed waveforms are the maximum currents across channels. The strongest stimulation pulses were delivered shortly after touch onset and offset. This resembled the natural temporal pattern of VPL spike rate in that almost all touch-responsive units recorded showed rapid adaptation.

In order to assess the spatiotemporal correlation between VPL rate and optimized ITMS, we estimated the cross-correlation between the two signals for each stimulating channel. In order to reduce the effect of suboptimal somatotopic coverage/representation on our analysis, the correlation was only computed for most accurate touch location in each animal. The cross-correlation function $R_{xy}(\tau)$ between two real signals x and y , and an unbiased estimate $\hat{R}_{xy}(\tau)$ using T samples

are defined as

$$R_{xy}(\tau) = E \{ x(t + \tau)y(t) \}$$

$$\hat{R}_{xy}(\tau) = \begin{cases} \frac{1}{T - |\tau|} \sum_{t=0}^{T-|\tau|-1} x(t + \tau)y(t), & \tau \geq 0 \\ \frac{1}{T - |\tau|} \hat{R}_{yx}(-\tau), & \tau < 0. \end{cases}$$

This formula was applied to estimate the function $\hat{R}_{\text{Rate,ITMS}}(\tau)$ for each touch pattern. Averages for each touch pattern, across all animals, are shown in figure 7(c) (top inset).

3.4. Touch parameter decoding

To assess the information content of virtual touch responses, we performed a set of classification experiments in which the touch condition (duration, location, amplitude) was predicted from multichannel peri-stimulus responses. This was first done separately for virtual and natural touch, to see how much information the neural responses in each modality provides about touch parameters. We then attempted to classify virtual touch responses under a single generalized classifier whose classification is based on the label of the nearest natural touch mean in a subspace optimized for both natural and virtual touch.

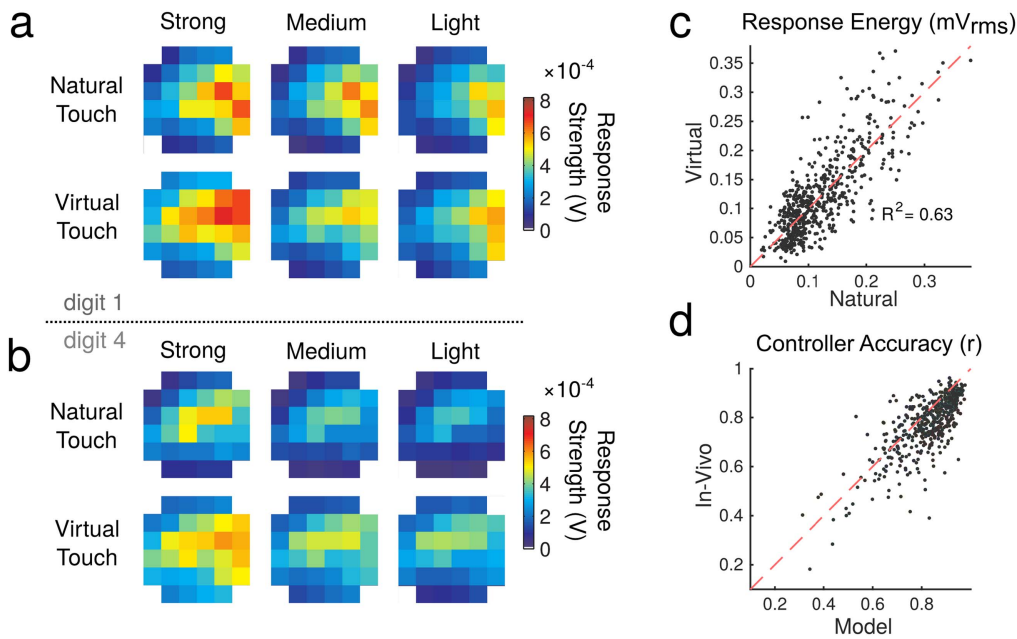


Figure 6. (a) and (b) Spatial response topologies natural and virtual touch. Each pixel corresponds with a recording electrode, and the color indicates the LFP strength, which we define as the most negative voltage in the response window (0–300 ms). (c) Channel-average amplitude comparison for natural touch and optimized ITMS *in vivo*. Each point represents a unique touch site/pattern combination. (d) A comparison between the modeled reproduction accuracy (correlation coefficient) of optimized ITMS and the actual accuracy obtained *in vivo*. Each point corresponds with a touch condition, and the touch site is color coded according to the touch site (see inset).

For natural and virtual touch, the individually trained classifiers could predict touch conditions with 56% and 61% accuracy, respectively. Given that the classification problems contained 30–54 classes, depending on the animal, this accuracy is quite high. The mutual information between the true class label and the estimated class label varied from 2.57 to 5.55 bits, with an average of close to 4 bits for both natural and virtual touch. The generalized classifier yielded classification rates of close to that of the individually trained classifiers (52% for natural touches, 54% for virtual touches). This means that using the combined responses to learn the LDA projection did not degrade the discriminative quality present in the individually-learned projections. Table 1 shows the classification performance and mutual information resulting from both the individually trained classifiers and the single generalized classifier. Table 2 shows the accuracy of classifying touch location when only considering trials of strong stimuli and short (150 ms) duration. In this case, the accuracy of decoding touch location alone was 90% for both natural and virtual touches. Chance levels in both tables reflect $1/(\text{number of classes})$, where the number of classes is either the number of touch conditions (table 1) or the number of sites (table 2). Using the average touch frequency for each animal, table 3 shows the information rate in bits s^{-1} between the true touch label (location, duration, and amplitude) and the estimated touch label. The virtual touches carry approximately the same information rate as the natural touches across the animals. The average information rate across animals is 5.2 bits s^{-1} for both natural and virtual touches trained separately, and the average information rate of virtual touches using the generalized classifier is 6.64 bits s^{-1} .

The joint classification rates for natural touch responses measured on the Utah array were on average 1.8-fold higher than those with the Michigan probe, and this was significant ($p = 0.024$, Wilcoxon rank sum test). As in the previously mentioned results for spatial reproduction accuracy, this can be explained by the greater somatotopic differentiability afforded by an array that spreads its channels horizontally across the surface of cortex rather than across layers. An analysis of conditional classification of touch pattern given the touch site revealed no significant difference ($p = 0.9$) in accuracy for natural responses. A smaller (1.3-fold higher), but significant difference in accuracy for virtual touch responses ($p = 0.024$) was observed.

Under the generalized classifier described above, virtual touch responses showed comparable levels of discriminability among different touch pressures, locations, and durations as natural touch responses. Figures 8(a) and (b) show classification rates based on two restricted subsets of trials. Figure 8(a) shows correct decoding rates for the subsets of data corresponding to light, medium, or strong touches of short (150 ms) duration. Figure 8(b) shows similar rates for both touch durations. In both of these classification experiments, stronger stimuli resulted in higher correct classification rates, which is not very surprising given that stronger touches were more accurately reproduced. The first column of figure 8(c) shows decoding performance when considering all types of trials, and the second column shows the overall accuracy when considering site-conditional subsets of data. The classification rates shown use the response in a window of 300 ms after touch onset for decoding. The classification rates across a range of windows sizes is shown in figure 8(d)

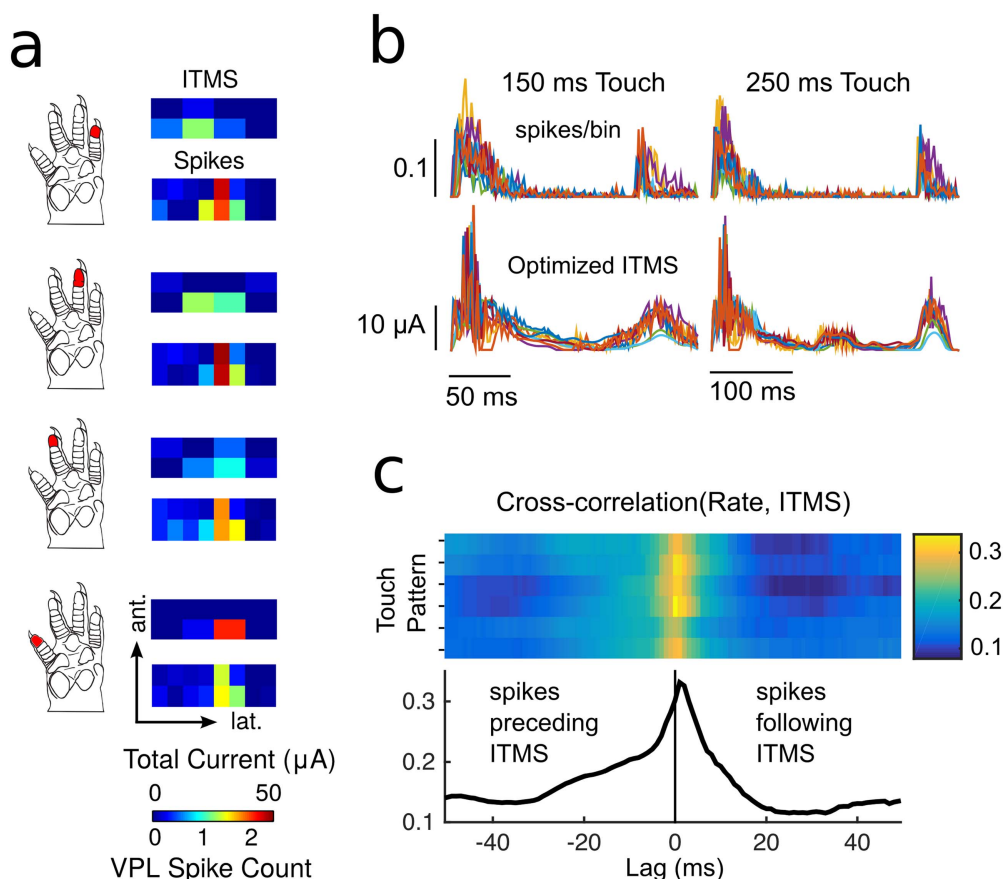


Figure 7. (a) Spatial distribution of electrode usage during optimized ITMS and VPL spiking activity during natural touch. For the four touch sites shown, the distribution of current injection on the VPL electrode array (top) and the accompanying average spike counts on the same electrodes (bottom) is shown. Spatially, the optimized current injection pattern coincides with the locations of responsive units. (b) The top row shows the post-onset MUA responses in VPL during natural touch and the corresponding optimized ITMS. Each trace represents the spike rate (units: spikes/bin) for a particular touch site. The two columns correspond with two different touch durations of 150 and 250 ms. Bottom inset: corresponding optimized ITMS, shown as the maximum current across stimulating channels. (c) Cross-correlation between VPL spike rate and optimized ITMS computed over all channels, touch sites, and animals. Top inset: cross-correlation for each touch pattern. Bottom inset: average across all patterns.

for decoding the touch location and amplitude given touches of fixed duration. It can be seen that for decoding touch site and amplitude, the classification rate reaches peak values at 15–20 ms after touch onset and remains high throughout the touch window, with a small increase shortly after touch offset.

4. Discussion

We have developed a neurophysiological approach to encoder design that optimizes the naturalness of downstream evoked responses. This provides a way to directly compute extremely detailed spatiotemporal microstimulation patterns that, according to a model of activation, are optimal for evoking desired natural responses. This study is, to our knowledge, the first *in vivo* demonstration of such a method.

This method, tested over a range of touch amplitudes and patterns, performed more accurately for short, strong touch patterns. This might suggest that a neuroprosthetic that uses this type of optimization would more effectively convey sensation of contact events or object slip rather than sustained

pressure. However, the result could have been due in part to the natural neural response exhibiting mostly onset-offset tuning in the neural subregions that were implanted. Generally, the strength of the neural readout was very predictive of controller accuracy, and had the responses shown more sustained-pressure tuning, the optimization might have conveyed this aspect of touch more informatively. Additionally, the spatial (somatotopic) reproduction accuracy with a readout array with greater horizontal coverage exceeded that of an array with good laminar but poor horizontal coverage. Since the same stimulating array was used in both cases, the difference in performance cannot be attributed to differences in controllability. Rather, the greater horizontal coverage puts more resolution in the dimensions that are most relevant for distinguishing somatotopically varying responses, and this increased observability is likely responsible for the increased accuracy. This suggests that a combination of horizontal and dorsoventral sampling could further improve overall control accuracy.

The optimized waveforms shared some notable characteristics: (1) for the physical contact area used, (1 mm²),

Table 1. Classification performance in terms of accuracy and mutual information for decoding touch parameters (touch site, duration, amplitude) from responses to natural and virtual touch. The mean and standard deviation of the rates across 8 Monte-Carlo divisions of data (2/3 train, 1/3 test) are shown. The individually trained classifiers used trial data from exclusively natural or virtual touch, while the dimensionality reduction for the joint classifier was trained with trials from both sets, and test examples are classified by selecting the nearest natural touch mean. For classification rate, the chance levels of prediction, which are dependent on the number of touch sites attempted on each animal, are shown in the second column. Chance levels: $1/(\text{number of classes})$. For mutual information, the entropy of the true touch labels is the upper bound and is shown in the second column.

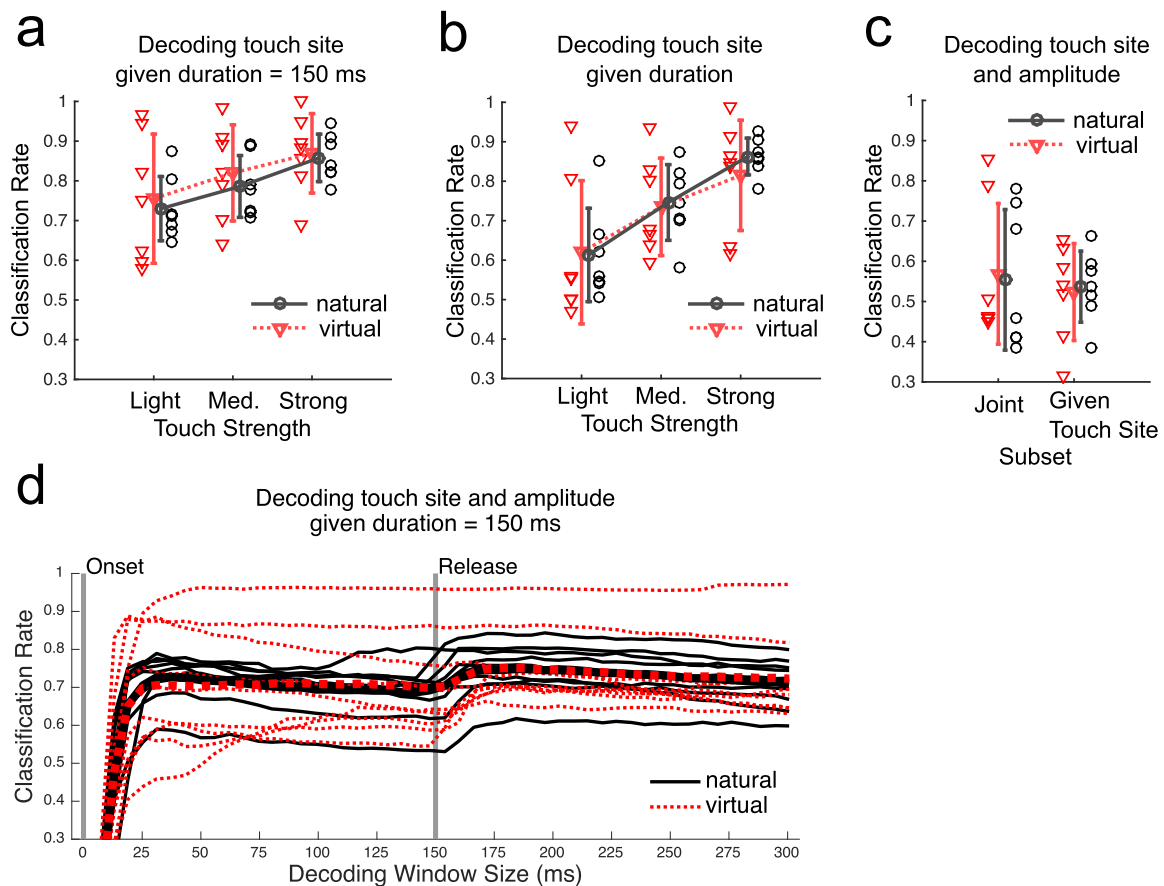
Animal		Trained separately				Generalized classifier			
		Natural		Virtual		Natural		Virtual	
		Mean	s.d.	Mean	s.d.	Mean	s.d.	Mean	s.d.
	Chance	Classification rate							
A	0.06	0.74	0.03	0.54	0.03	0.68	0.03	0.47	0.03
B	0.02	0.84	0.02	0.82	0.01	0.79	0.02	0.79	0.01
C	0.02	0.81	0.02	0.66	0.02	0.75	0.02	0.51	0.02
D	0.02	0.40	0.02	0.95	0.01	0.38	0.03	0.85	0.02
E	0.02	0.44	0.02	0.54	0.03	0.41	0.02	0.46	0.01
F	0.02	0.47	0.02	0.51	0.02	0.45	0.02	0.46	0.02
G	0.02	0.46	0.03	0.45	0.02	0.41	0.02	0.44	0.02
H	0.02	0.52	0.02	0.45	0.02	0.47	0.02	0.37	0.02
I	0.02	0.37	0.02	0.54	0.02	0.36	0.02	0.51	0.02
Average		0.56	0.18	0.61	0.17	0.52	0.17	0.54	0.16
	Entro- py	Mutual information (bits)							
A	4.17	3.21	0.10	2.57	0.06	2.98	0.09	3.26	0.10
B	5.58	4.98	0.06	5.01	0.04	4.81	0.06	5.31	0.04
C	5.75	5.00	0.06	4.59	0.08	4.76	0.06	5.10	0.04
D	5.74	3.67	0.06	5.55	0.03	3.59	0.07	5.49	0.04
E	5.73	3.79	0.06	3.84	0.08	3.73	0.05	4.87	0.07
F	5.73	3.88	0.07	3.58	0.09	3.83	0.08	4.87	0.06
G	5.73	3.81	0.09	3.40	0.09	3.63	0.05	4.69	0.05
H	5.73	4.11	0.07	3.68	0.07	3.90	0.07	4.40	0.08
I	5.73	3.34	0.07	3.74	0.08	3.40	0.06	4.85	0.06
Average		3.98	0.64	4.00	0.90	3.85	0.60	4.76	0.65

Table 2. Decoding touch site given a short, strong touch. Correct classification rates here reflect the discriminability endowed by spatial variation in natural and/or virtual stimuli. Similar to table 1, the chance level is $1/(\text{number of classes})$. The means and standard deviations shown are with respect to 8 Monte-Carlo trial divisions (2/3 train, 1/3 test).

Animal		Trained separately				Generalized classifier			
		Natural		Virtual		Natural		Virtual	
		Mean	s.d.	Mean	s.d.	Mean	s.d.	Mean	s.d.
A	0.33	0.98	0.02	0.93	0.03	0.95	0.03	0.89	0.04
B	0.12	0.95	0.03	0.96	0.03	0.89	0.02	0.95	0.03
C	0.11	0.90	0.03	0.81	0.04	0.80	0.05	0.66	0.04
D	0.11	0.93	0.03	1.00	0.01	0.91	0.03	1.00	0.01
E	0.11	0.86	0.05	0.85	0.04	0.77	0.04	0.84	0.04
F	0.11	0.91	0.03	0.90	0.03	0.81	0.05	0.79	0.04
G	0.11	0.90	0.04	0.92	0.03	0.83	0.03	0.89	0.03
H	0.11	0.91	0.03	0.86	0.04	0.89	0.04	0.84	0.06
I	0.11	0.77	0.05	0.86	0.03	0.72	0.05	0.86	0.05
Average		0.90	0.06	0.90	0.06	0.84	0.07	0.86	0.10

Table 3. Average information rate (bits s^{-1}) for decoding touch parameters (touch site, duration, amplitude) from responses to natural and virtual touch.

Animal	Touch rate	Trained separately		Generalized classifier	
		Natural	Virtual	Natural	Virtual
A	2.45 Hz	7.85	6.29	7.28	7.98
B	0.91 Hz	4.52	4.55	4.37	4.82
C	0.91 Hz	4.57	4.19	4.35	4.66
D	1.43 Hz	5.26	7.96	5.15	7.88
E	1.45 Hz	5.49	5.56	5.39	7.04
F	1.46 Hz	5.66	5.23	5.60	7.10
G	1.47 Hz	5.58	4.99	5.32	6.88
H	1.45 Hz	5.97	5.34	5.67	6.39
Average		5.52	5.50	5.34	6.64

**Figure 8.** Performance in correctly inferring touch parameters from natural (black) and virtual (red) touch responses. Decoding is performed using a generalized classifier—one whose LDA projection is trained using responses from both natural and virtual touch. For prediction, new examples are first projected and then classified based on nearest natural touch mean. Classification rate is calculated as the (number of correctly classified trials)/(total number of trials) computed over 8 Monte-Carlo selections of training/test data (2/3 train, 1/3 test). Data points represent animals. Error bars show ± 1 s.d. across animals. (a) Decoding rates when only considering trials that were of short (150 ms) duration and were light, medium, or strong. (b) Classification rate for decoding duration and touch site when the touch strength is known. (c) Classification rate for all trials in the first column. The second column shows the accuracy in decoding amplitude and duration when the touch site is known. (d) Classification rate across different window sizes. Thin lines represent animals. Thick lines show the average across animals.

most of the optimized ITMS patterns used 1–3 electrode configurations over the course of 300 ms following touch onset, and these configurations were usually adjacent to each other. (2) Temporally, the optimized ITMS consisted of a brief burst of pulses beginning 5–7 ms after touch onset and a

secondary burst of pulses shortly after touch release. In between these two bursts, the amount of charge injected in the holding period was negligible. (3) The stimulation amplitudes (with stereotyped biphasic pulses with 200 μs per phase) remained largely below 30 μA per phase, although in this

work we did not aggressively constrain the maximum input current, but rather used the range of current that was applied during the probing microstimulation sequence. Future work might measure the relative effectiveness of highly constrained input currents, which might prove useful for smaller electrode surface areas.

4.1. Information transfer

Ideally, a sensory prosthetic encoder should accomplish two related goals: high information transfer and natural neural activation. Much of the work so far in somatosensory prosthetics has focused on information transfer, measured by perceptual discriminability. Our work, in contrast, focused on the naturalness of neural activation, and we found, through a series of classification experiments, that optimized responses showed the same amount of discriminability of touch parameters as natural responses. Single-trial LFP responses to both natural and virtual touch were quite discriminable with an average accuracy of 56% and 61%, respectively, in decoding the joint parameters corresponding to site, amplitude, and duration. The mutual information between the true touch type and inferred touch was 4 bits in both cases. Given a strong short touch, the touch sites were predicted with 90% accuracy for both virtual and natural touch. By varying the size of the window used for decoding, it is evident that the touch location and amplitude can be reliably decoded using only the first 25 ms of the touch response for touches of fixed duration. Thus, most of the information on amplitude and location is present within this initial window, and a similar window following release (figure 8(d)). This corroborates with results that found that the time points 12–14 ms after touch onset are the most important for touch location decoding accuracy [32].

Virtual touch responses to different touch conditions were not only discriminable from each other, but were also separable along the same boundaries as natural touch responses. When virtual touch responses were classified according to their similarity to corresponding natural touch means, the resulting classification rate (52%) matched that of natural touch responses. This demonstrates that our microstimulation evoked responses were not only informative by exhibiting high parameter discriminability, but were natural by showing preferential similarity to the natural touch counterparts. Similarly, when the accuracy of virtual touch responses were measured using non-matched touch labels (figure 5(e)), the scores were significantly worse.

Information transfer is the motivation behind many discrimination studies wherein isolated patterning parameters are studied psychometrically. In another approach by [44], information rate is specifically optimized by posing microstimulation patterning as a channel coding problem, designing a transduction filter that maximizes the mutual information between external stimuli and the neural response (in a neural model of the thalamocortical system). This framework could, however, lead to encoders with high information transfer but very unnatural spatiotemporal mappings. Our method, in contrast, does not optimize for information transfer explicitly,

but produces responses that are discriminable if such information was evident in the natural responses. In fact, the virtual touches carried approximately the same information rate as the natural touches; for both natural and virtual touches, the average information rate across animals was 5.5 bits s^{-1} . A future study might explore explicitly combining information transfer and naturalness into a joint objective that could be optimized.

Features of touch onsets, such as their amplitude and spatial location, could be discriminated from virtual touches with the same latency, on average, as their natural touch counterparts. As shown in figure 8(d), this latency was 15–20 ms after touch onset, which corresponds with previously reported peak spiking latencies in the forelimb area of somatosensory cortex [45]. Virtual touch stimuli therefore not only provide naturalistic levels of information on touch parameters, they do so with the same timing, as would be expected for a biomimetic sensory prosthesis.

The discriminability of natural touch responses depended significantly on the geometry of the electrode array. Specifically, the Utah array, which distributes its channels horizontally in a 2D grid across cortical surface, led to more accurate classification than the Michigan array, whose contacts span one horizontal and one laminar axis. The improvement could be due to the Utah array's less ambiguous recording of somatotopic differences between touch locations, an example of which is shown in figure 6. This is supported by a similar analysis of classification of touch pattern conditioned on touch site that revealed no significant difference in accuracy between the two arrays. While responses from the Michigan probe still contain information on touch location (see table 2), they are limited by the fact that the probe can only resolve horizontal details along a single axis. Interestingly, in an isolated case with a Michigan probe (rat D in table 1), the virtual touch decoding accuracy was much higher than that of natural touch. However, for short strong touches in table 2, the classification rates were comparable. This can be explained by the variability of virtual touch responses being much lower than the natural response variability—a pattern observed only in this particular animal. Lower variability would lead to fewer misclassifications of virtual responses as long as their relationships to the nearest natural touch means were accurate. In the comparatively easier problem of classifying touch sites alone, this variability played less of a role, and the natural/virtual discrepancy was greatly decreased.

4.2. Emulating the neural code in the stimulation area

Some groups have shown that an encoder that mimics the natural spiking activity of an implanted region can be somewhat effective in terms of neural or psychophysical readout, either by playback of recorded SUA [12], or forward point process modeling [46, 47]. Our work differs in the sense that the stimulation does not explicitly follow the spike rate in the stimulated brain region, but rather is optimized to evoke downstream responses similar to natural touch.

In experiments, our method performed favorably compared to rate-based stimulation, showing that the optimization compensates for some spatial and dynamical effects of microstimulation, which has been shown to activate neural elements that are not directly measurable from single-unit recordings, such as fibers of passage. With each pulse, it is difficult to ascertain how many cells were activated, and it has been shown that the projection fields of stimulating pulses—the somatotopic topology of their downstream responses—are offset from the receptive fields on units recorded on the same electrodes [48]. Therefore, while rate-based microstimulation does indeed resemble the neural code at the stimulated region, the spatial topology of their responses as well as the temporal trajectories could be far from the desired natural activation. Interestingly, in post-hoc analyses, we showed that the optimized control inputs qualitatively resembled the background-subtracted VPL PSTHs in the sense that both signals exhibited rapidly-adapting tuning and involved a somewhat overlapping set of electrode channels. However, quantitative analyses showed that this was a weak correlation, and the two signals were quite different spatially and temporally.

4.3. Linearity of responses

Neural responses to single-channel thalamic and cortical microstimulation have been shown to exhibit nonlinear effects such as paired-pulse facilitation and attenuation [13, 31, 46] for pulses within 200 ms of each other. In our models, the only nonlinearity was an input threshold, so these temporal interaction effects were not modeled. It is possible that some inaccuracies were due in part to unmodeled nonlinear effects, but if this were true, the theoretical accuracy under our model would greatly exceed experimental accuracy. In contrast, we observed that the theoretical accuracy was only 5.8% more accurate than that observed experimentally, suggesting that the model was not the primary source of error. Although more accurate models of activation could be trained, [27, 46, 49–51], they are more computationally involved to control—often without assurance of optimality.

4.4. Obtaining neural templates

Subject-specific neural responses to natural stimuli would not be available in a somatosensory prosthetic setting, since the target population for such a device would, by definition, lack intact somatosensory representation. This sort of problem is certainly not unique to sensory neural prostheses. Most of the work on motor brain-machine interfaces [52–54] uses intact limb kinematics in non-human primates to initialize models that map neural firing to limb kinematics and/or force. More recently, studies on hippocampal prosthetics such as those conducted by Berger *et al* require full observation of neural firing from input and output populations to train a mapping.

Nevertheless, fully observed experiments such as these can help identify patterns that can be generalized across subjects—and do so with a throughput not available with purely psychophysical methods. Although it remains to be

seen how effectively a full spatiotemporal encoder can generalize across subjects, it is already known that several temporal features can generalize across somatotopic locations and subjects. In [4] it was shown that for a simple encoder consisting of a static nonlinearity, the best parameter values were remarkably similar across subjects and electrode sites. In addition, flutter frequency percepts have been shown to generalize across subjects, provided that the stimulus location contains rapidly adapting cells [2]. These studies suggest that, given spatiotemporal patterns that produce naturalistic responses in one subject, a simple somatotopic realignment could sufficiently restore near-naturalistic responses in other subjects.

4.5. Translational applicability

Techniques for optimally controlling artificial input to neural systems to restore realistic sensory responses are only beginning to be validated *in vivo*, and to our knowledge, our study is the first application to microstimulation of somatosensory circuits. Although the same approach applied to the non-human primate would provide data that could be more relevant for neuroprosthetics, it would likely be slower and more costly to iterate. The rat model, which exhibits several fundamental similarities to the primate somatosensory system [20], provides a high-throughput testbed for validating these nascent techniques. Furthermore, the models and controllers in this work make only broad assumptions about the underlying physiology—nothing that precludes immediate extension to the primate system, where species-specific refinements could be introduced.

In the present work, ITMS patterns were optimized separately for each touch condition, but an important goal for future work would be to build monolithic encoders that convert arbitrary spatiotemporal touch patterns to microstimulation using a single set of parameters. These would likely be nonlinear, fairly complex spatiotemporal models, perhaps involving recurrent neural networks or a similarly rich class of multi-input multi-output filters. These models are beyond the scope of this work, but we note that separate optimization, like the kind used in this study, provides an upper bound on control accuracy since it is not constrained by having to account for the full range of touch conditions. Separately optimized microstimulation patterns could also be used as training data for a monolithic encoder, where spatiotemporal touch input can be functionally related to (pre)optimized microstimulation, obtained using our methods. This converts the potentially difficult problem of learning a nonlinear controller into the comparatively easier task of identifying a dynamical system [27, 49]. To emulate natural somatosensation, future studies should explore encoding and decoding touches at multiple skin locations and of overlapping duration. Most generally, trials could consist of randomly distributed patterns applied spatiotemporally across the skin surface, similar in principle to those used in visual [55, 56] and vibrissal [57] studies.

For human application, we envision this potentially high dimensional optimization being performed through exhaustive experiments in animal models, and simpler, lower dimensional calibrations being subsequently fine-tuned for human patients. For example, since the limb representation of VPL and S1 are somatotopically organized, it is possible that cross-hemispheric, cross-subject or cross-species generalization could be largely accomplished through simple intensity scaling and/or spatial remapping. These calibrations could also be optimized under a reinforcement learning framework [58] in which user-generated evaluative feedback could drive fine adjustments to parameters over time.

Acknowledgments

This work was supported by the DARPA REPAIR project N66001-10-C-2008 and the New York State Spinal Cord Injury Research Board Awards 73525 (C30600GG) and 73689 (DOH01C30838GG3450000).

References

- [1] Hsiao S S, Fettiplace M and Darbandi B 2011 *Prog. Brain Res.* **192** 69–81
- [2] Romo R, Hernández A, Zainos A, Brody C D and Lemus L 2000 *Neuron* **26** 273–8
- [3] Berg J A et al 2013 *IEEE Trans. Neural Syst. Rehabil. Eng.* **21** 500–7
- [4] Tabot G A, Dammann J F, Berg J A, Tenore F V, Boback J L, Vogelstein R J and Bensmaia S J 2013 *Proc. Natl Acad. Sci. USA* **110** 18279–84
- [5] O'Doherty J E, Lebedev M A, Hanson T L, Fitzsimmons N A and Nicolelis M A L 2009 *Frontiers Integrative Neurosci.* **3** 20
- [6] Fitzsimmons N A, Drake W, Hanson T L, Lebedev M A and Nicolelis M A L 2007 *J. Neurosci.* **27** 5593–602
- [7] London B M, Jordan L R, Jackson C R and Miller L E 2008 *IEEE Trans. Neural Syst. Rehabil. Eng.* **16** 32–6
- [8] O'Doherty J E, Lebedev M A, Li Z and Nicolelis M A L 2012 *IEEE Trans. Neural Syst. Rehabil. Eng.* **20** 85–93
- [9] Heming E, Sanden A and Kiss Z H T 2010 *J. Neural Eng.* **7** 064001
- [10] Patel S, Ohara S, Dougherty P M, Gracely R H and Lenz F A 2006 *J. Neurophysiol.* **95** 646–59
- [11] Berger T W, Hampson R E, Song D, Goonawardena A, Marmarelis V Z and Deadwyler S A 2011 *J. Neural Eng.* **8** 046017
- [12] Weber D J, London B M, Hokanson J A, Ayers C A, Gaunt R A, Torres R R, Zaaimi B and Miller L E 2011 *IEEE Trans. Neural Syst. Rehabil. Eng.* **19** 501–13
- [13] Butovas S and Schwarz C 2003 *J. Neurophysiol.* **90** 3024–39
- [14] Stoney S D, Thompson W D and Asanuma H 1968 *J. Neurophysiol.* **31** 659–69
- [15] Histed M H, Bonin V and Reid R C 2009 *Neuron* **63** 508–22
- [16] Pillow J W, Ahmadian Y and Paninski L 2011 *Neural Comput.* **23** 1–45
- [17] Paninski L, Pillow J and Lewi J 2007 *Prog. Brain Res.* **165** 493–507
- [18] Ahmadian Y, Packer A M, Yuste R and Paninski L 2011 *J. Neurophysiol.* **106** 1038–53
- [19] Li L, Park I M, Brockmeier A, Chen B, Seth S, Francis J T, Sanchez J C and Principe J C 2013 *IEEE Trans. Neural Syst. Rehabil. Eng.* **21** 532–43
- [20] Francis J T, Xu S and Chapin J K 2008 *J. Neurophysiol.* **99** 2291–304
- [21] Paxinos G and Watson C 2006 *The Rat Brain in Stereotaxic Coordinates* (New York: Academic)
- [22] Foffani G, Chapin J K and Moxon K A 2008 *J. Neurophysiol.* **100** 268–80
- [23] Gaunt R A, Hokanson J A and Weber D J 2009 *J. Neural Eng.* **6** 055009
- [24] Ranck J B 1975 *Brain Res.* **98** 417–40
- [25] Cogan S F 2008 *Annu. Rev. Biomed. Eng.* **10** 275–309
- [26] Merrill D R, Bikson M and Jefferys J G R 2005 *J. Neurosci. Methods* **141** 171–98
- [27] Millard D C, Wang Q, Gollnick C A and Stanley G B 2013 *J. Neural Eng.* **10** 066011
- [28] Wu Y T and Scabassi R J 1997 *IEEE Trans. Syst. Man Cybern. A* **27** 222–34
- [29] Scabassi R J, Hinman C L, Kroin J S and Risch H A 1985 *Electroencephalogr. Clin. Neurophysiol.* **60** 444–54
- [30] Berger T W, Eriksson J L, Ciarolla D A and Scabassi R J 1988 *J. Neurophysiol.* **60** 1095–109
- [31] Castro-Alamancos M A and Connors B W 1996 *J. Neurosci.* **16** 2767–79
- [32] Brockmeier A J, Choi J S, Kriminger E G, Francis J T and Principe J C 2014 *Neural Comput.* **26** 1080–107
- [33] Li L, Brockmeier A J, Choi J S, Francis J T, Sanchez J C and Principe J C 2014 *Comput. Intell. Neurosci.* **2014** 870160
- [34] Markowitz D A, Wong Y T, Gray C M and Pesaran B 2011 *J. Neurosci.* **31** 18412–22
- [35] Stavisky S D, Kao J C, Nuyujukian P, Ryu S I and Shenoy K V 2015 *J. Neural Eng.* **12** 036009
- [36] Flint R D, Lindberg E W, Jordan L R, Miller L E and Slutsky M W 2012 *J. Neural Eng.* **9** 046006
- [37] Van Overschee P and De Moor B 2011 *Subspace Identification for Linear Systems Theory—Implementation—Applications* (Berlin: Springer)
- [38] Ljung L 1999 *System Identification* (Upper Saddle River, NJ: Prentice Hall)
- [39] Boyd S and Vandenberghe L 2004 *Convex Optimization* (New York: Cambridge University Press)
- [40] Wang Y and Boyd S 2010 *IEEE Trans. Control Syst. Technol.* **18** 267–78
- [41] Koivuniemi A S and Otto K J 2011 *IEEE Trans. Neural Syst. Rehabil. Eng.* **19** 468–76
- [42] Li W and Todorov E 2004 Iterative linear quadratic regulator design for nonlinear biological movement systems *Proc. Int. Conf. on Informatics in Control, Automation and Robotics* pp 222–9
- [43] Fukunaga K 1990 *Introduction to Statistical Pattern Recognition* 2nd edn (San Diego, CA: Academic)
- [44] Daly J, Liu J, Aghagolzadeh M and Oweiss K 2012 *J. Neural Eng.* **9** 065004
- [45] Tutunculer B, Foffani G, Himes B T and Moxon K A 2006 *Cerebral Cortex* **16** 791–810
- [46] Song D, Chan R H M, Marmarelis V Z, Hampson R E, Deadwyler S A and Berger T W 2009 *Neural Netw.* **22** 1340–51
- [47] Nirenberg S and Pandarinath C 2012 *Proc. Natl Acad. Sci. USA* **109** 15012–7
- [48] Wang Q, Millard D C, Zheng H J V and Stanley G B 2012 *J. Neural Eng.* **9** 026008
- [49] Marmarelis P V Z 2004 *Nonlinear Dynamic Modeling of Physiological Systems* (New York, Piscataway, NJ: Wiley-IEEE)
- [50] Wu M C K, David S V and Gallant J L 2006 *Annu. Rev. Neurosci.* **29** 477–505
- [51] Marmarelis V Z 1993 *Ann. Biomed. Eng.* **21** 573–89

- [52] Carmena J M, Lebedev M A, Crist R E, O'Doherty J E, Santucci D M, Dimitrov D F, Patil P G, Henriquez C S and Nicolelis M A L 2003 *PLoS Biol.* **1** e42
- [53] Serruya M D, Hatsopoulos N G, Paninski L, Fellows M R and Donoghue J P 2002 *Nature* **416** 141–2
- [54] Taylor D M, Tillery S I H and Schwartz A B 2002 *Science* **296** 1829–32
- [55] Chichilnisky E J 2001 *Network* **12** 199–213
- [56] Sharpee T O 2013 *Annu. Rev. Neurosci.* **36** 103–20
- [57] Ramirez A, Pnevmatikakis E A, Merel J, Paninski L, Miller K D and Bruno R M 2014 *Nat. Neurosci.* **17** 866–75
- [58] Brockmeier A J, Choi J S, Distasio M M, Francis J T and Príncipe J C 2011 Optimizing microstimulation using a reinforcement learning framework 2011 *Annual Int. Conf. of the IEEE Engineering in Medicine and Biology Society, EMBC* (Piscataway, NJ: IEEE) pp 1069–72